

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: June 10, 2013

G. Bertrand, Ed.
E. Stephan
France Telecom - Orange
R. Peterkofsky
Skytide, Inc.
F. Le Faucheur
Cisco Systems
P. Grochocki
Orange Polska
December 7, 2012

CDNI Logging Interface
draft-ietf-cdni-logging-00

Abstract

This memo specifies the Logging interface between a downstream CDN (dCDN) and an upstream CDN (uCDN) that are interconnected as per the CDN Interconnection (CDNI) framework. First, it describes a reference model for CDNI logging. Then, it specifies the actual protocol for CDNI logging information exchange covering the information elements as well as the transport of those.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on April 25, 2013.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents

(<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1.	Introduction	5
1.1.	Terminology	6
1.2.	Abbreviations	8
2.	CDNI Logging Reference Model	8
2.1.	CDNI Logging interactions	8
2.2.	Overall Logging Chain	13
2.2.1.	Logging Generation and During-Generation Aggregation	15
2.2.2.	Logging Collection	15
2.2.3.	Logging Filtering	16
2.2.4.	Logging Rectification and Post-Generation Aggregation	16
2.2.5.	Log-Consuming Applications	17
2.2.5.1.	Maintenance/Debugging	17
2.2.5.2.	Accounting	17
2.2.5.3.	Analytics and Reporting	18
2.2.5.4.	Security	18
2.2.5.5.	Legal Logging Duties	18
2.2.5.6.	Notions common to multiple Log Consuming Applications	18
3.	CDNI Logging Information Structure and Transport	20
4.	CDNI Logging Fields	22
4.1.	Generic Fields	22
4.1.1.	Semantics of Generic CDNI Logging Fields	22
4.1.2.	Syntax of Generic CDNI Logging Fields	24
4.2.	Logging Fields for Content Delivery	25
4.2.1.	Semantics for Delivery CDNI Logging Fields	25
4.2.2.	Syntax for Delivery CDNI Logging Fields	26
4.3.	Logging Fields for Content Acquisition	26
4.3.1.	Semantics for Acquisition CDNI Logging Fields	27
4.3.2.	Syntax for Acquisition CDNI Logging Fields	27
4.4.	Logging Fields for Control	27
4.5.	Logging Fields for Other Operations	27
5.	CDNI Logging Records	28
5.1.	Content Delivery	28
5.2.	Content Acquisition	29
5.2.1.	Logging Records Provided by dCDN to uCDN	29
5.2.2.	Logging Records Provided by uCDN to dCDN	29
5.3.	Content Invalidation and Purging	30
5.4.	Logging Extensibility	30
6.	CDNI Logging File Format	30
6.1.	Logging Files	31
6.2.	File Format	31
6.2.1.	Headers	31
6.2.2.	Body (Logging Records) Format	32
6.2.3.	Footer Format	33

7.	CDNI Logging File Transport Protocol	33
8.	Logging Control	33
9.	Open Issues	34
10.	IANA Considerations	35
11.	Security Considerations	35
11.1.	Privacy	35
11.2.	Non Repudiation	35
12.	Acknowledgments	35
13.	References	35
13.1.	Normative References	35
13.2.	Informative References	36
Appendix A.	Examples Log Format	37
A.1.	W3C Common Log File (CLF) Format	37
A.2.	W3C Extended Log File (ELF) Format	38
A.3.	National Center for Supercomputing Applications (NCSA) Common Log Format	39
A.4.	NCSA Combined Log Format	39
A.5.	NCSA Separate Log Format	39
A.6.	Squid 2.0 Native Log Format for Access Logs	40
Appendix B.	Requirements	41
B.1.	Additional Requirements	41
B.2.	Compliance with Requirements draft	42
Appendix C.	CDNI WG's position on candidate protocols for Logging Transport	42
C.1.	CDNI WG's position on Syslog	42
C.2.	CDNI WG's position on SNMP	42
Authors'	Addresses	42

1. Introduction

This memo specifies the Logging interface between a downstream CDN (dCDN) and an upstream CDN (uCDN). First, it describes a reference model for CDNI logging. Then, it specifies the actual protocol for CDNI logging information exchange covering the information elements as well as the transport of those.

The reader should be familiar with the work of the CDNI WG:

- o CDNI problem statement [[RFC6707](#)] and framework [[I-D.ietf-cdni-framework](#)] identify a Logging interface,
- o Section 7 of [[I-D.ietf-cdni-requirements](#)] specifies a set of requirements for Logging,
- o [[I-D.ietf-cdni-use-cases](#)] outlines real world use-cases for interconnecting CDNs. These use cases require the exchange of Logging information between the dCDN and the uCDN.

As stated in [[RFC6707](#)], "the CDNI Logging interface enables details of logs or events to be exchanged between interconnected CDNs".

The present document describes:

- o The CDNI Logging reference model ([Section 2](#)),
- o The CDNI Logging information structure and Transport ([Section 3](#)),
- o The CDNI Logging Fields ([Section 4](#)),
- o The CDNI Logging Records ([Section 5](#)),
- o The CDNI Logging File format ([Section 6](#)),
- o The CDNI Logging File Transport Protocol ([Section 7](#)),
- o and, finally, the description of the CDNI Logging Control that is to be supported by the CDNI Control Interface [Section 8](#).

In the Appendices, the document provides:

- o A list of identified requirements (Appendix B.1), which should be considered for inclusion in [[I-D.ietf-cdni-requirements](#)],

1.1.1. Terminology

In this document, the first letter of each CDNI-specific term is capitalized. We adopt the terminology described in [[RFC6707](#)] and [[I-D.ietf-cdni-framework](#)], and extend it with the additional terms defined below.

For clarity, we use the word "Log" only for referring to internal CDN logs and we use the word "Logging" for any inter-CDN information exchange and processing operations related to the CDNI Logging interface. Log and Logging formats may be different.

Log: CDN internal information collection and processing operations.

Logging: Inter-CDN information exchange and processing operations.

CDNI Logging Field: an atomic element of information that can be included in a CDNI Logging Record. The time an event/task started, the IP address of an End user to whom content was delivered, and the URI of the content delivered are examples of CDNI logging fields.

CDNI Logging Record: an information record providing information about a specific event. This comprises a collection of CDNI Logging Fields.

Separator Character: a specific character used to enable the parsing of Logging Records. This character separates the Logging Fields that compose a Logging Record.

Logging File: a file containing Logging Records and additional information for easing the processing of the Logging Records.

CDN Reporting: the process of providing the relevant information that will be used to create a formatted content delivery report provided to the CSP in deferred time. Such information typically includes aggregated data that can cover a large period of time (e.g., from hours to several months). Uses of Reporting include the collection of charging data related to CDN services and the computation of Key Performance Indicators (KPIs).

CDN Monitoring: the process of providing content delivery information in real-time. Monitoring typically includes data in real time to provide visibility of the deliveries in progress, for service operation purposes. It presents a view of the global health of the services as well as information on usage and performance, for network services supervision and operation management. In particular, monitoring data can be used to generate alarms.

End-User experience management: study of Logging data using statistical analysis to discover, understand, and predict user behavior patterns.

Class-of-requests: A Class-of-requests identifies a set of content Requests, related to a specific CSP, received from clients in a given footprint and sharing common properties. These properties include:

- o Any header, URL parameter, query parameter of an HTTP (or RTMP) content request
- o Any header, or sub-domain of the FQDN of a DNS lookup request

Examples:

- o Class-of-Requests = all the requests that include the HTTP header "User-Agent: Mozilla/5.0" related to CSP "http://*.cdn.example.com" from AS3215
- o Class-of-Requests = all the DNS requests from anywhere and related to CSP "cdn*.example.com"

Delivery Service: A Delivery Service is defined by a set of Class-of-Requests and a list of parameters that apply to all these Class-of-Requests (logging format, delivery quality/capabilities requirements...)

Service Agreement: A service agreement is defined by a uCDN identifier, a dCDN identifier, a set of Delivery Services and a list of parameters that apply to the Service Agreement.

Once a Service Agreement is agreed between the administrative entities managing the CDNs to be interconnected, the upstream CDN and the downstream CDN of the CDNI interconnection must be configured according to this agreed Service Agreement. For instance, a given uCDN (uCDN1) may request a given dCDN (dCDN1) to configure one Delivery Service for handling requests for HTTP Adaptive streaming videos delegated by uCDN1 and related to a specific CSP (CSP1) and another one for handling requests for static pictures delegated by uCDN1 and related to CSP1. These Delivery services would belong to the Service Agreement between uCDN1 and dCDN1 for CSP1. In this simple example, uCDN1 may request dCDN1 to include Delivery Service information in its CDNI Logging, to help uCDN1 to provide relevant reports to CSP1.

1.2. Abbreviations

- o API: Application Programming Interface
- o CCID: Content Collection Identifier
- o CDN: Content Delivery Network
- o CDNP: Content Delivery Network Provider
- o CoDR: Content Delivery Record
- o CSP: Content Service Provider
- o DASH: Dynamic Adaptive Streaming over HTTP
- o dCDN: downstream CDN
- o FTP: File Transfer Protocol
- o HAS: HTTP Adaptive Streaming
- o KPI: Key Performance Indicator
- o PVR: Personal Video Recorder
- o SID: Session Identifier
- o SFTP: SSH File Transfer Protocol
- o SNMP: Simple Network Management Protocol
- o uCDN: upstream CDN

2. CDNI Logging Reference Model

2.1. CDNI Logging interactions

The CDNI logging reference model between a given uCDN and a given dCDN involves the following interactions:

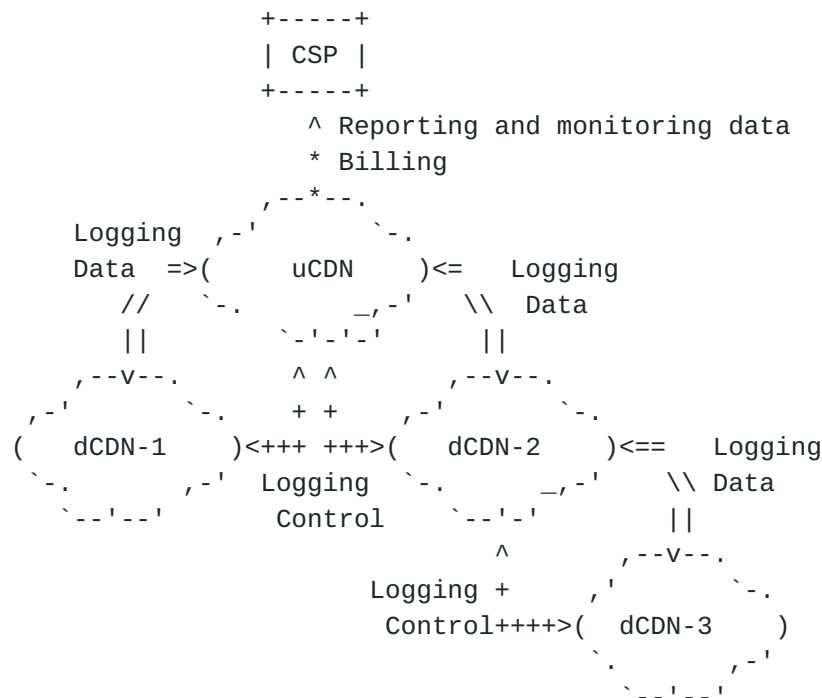
- o control by the uCDN of the logging to be performed by the dCDN (e.g. control of which logging fields are to be communicated to the uCDN for a given task performed by the dCDN, control of which types of events are to be logged). This is supported by the CDNI Control interface.

- o generation and collection by the dCDN of logging information related to the completion of any task performed by the dCDN on behalf of the uCDN (e.g. delivery of the content to an end user) or related to events happening in the dCDN that are relevant to the uCDN (e.g. failures or unavailability in dCDN). This takes place within the dCDN and does not directly involve CDNI interfaces.
- o communication by the dCDN to the uCDN of the logging information collected by the dCDN relevant to the uCDN. This is supported by the CDNI Logging interface. For example, the uCDN may use this logging information to charge the CSP, to perform analytics and monitoring for operational reasons, to provide analytics and monitoring views on its content delivery to the CSP, or to perform troubleshooting.
- o control by the dCDN of the logging to be performed by the uCDN on behalf of the dCDN. This is supported by the CDNI Control interface.
- o generation and collection by the uCDN of logging information related to the completion of any task performed by the uCDN on behalf of the dCDN (e.g. serving of content by uCDN to dCDN for acquisition purposes by dCDN) or related to events happening in the uCDN that are relevant to the dCDN. This takes place within the uCDN and does not directly involve CDNI interfaces.
- o communication by the uCDN to the dCDN of the logging information collected by the uCDN relevant to the dCDN. This is supported by the CDNI Logging interface. For example, the dCDN may use this logging information for security auditing or content acquisition troubleshooting.

Figure 1 provides an example of CDNI Logging interactions in a particular scenario where 4 CDNs are involved in the delivery of content from a given CSP: the uCDN has a CDNI interconnection with dCDN1 and dCDN2. In turn, dCDN2 has a CDNI interconnection with dCDN3. uCDN, dCDN1, dCDN2 and dCDN3 deliver content for the CSP. In this example, the CDNI Logging interface enables the uCDN to obtain logging information from all the dCDNs involved in the delivery. In the example, uCDN uses the Logging data:

- o to analyze the performance of the delivery operated by the dCDNs and to adjust its operations (e.g., request routing) as appropriate
- o to provide reporting (non-real time) and monitoring (real time) information to CSP.

For instance, uCDN merges Logging data, extracts relevant KPIs, and presents a formatted report to CSP, in addition to a bill for the content delivered by uCDN itself or its dCDNs on his behalf. uCDN may also provide Logging data as raw log files to CSP, so that CSP can use its own Logging analysis tools.



```
<====> CDNI Logging Interface
<++++> CDNI Control Interface
***> outside the scope of CDNI
```

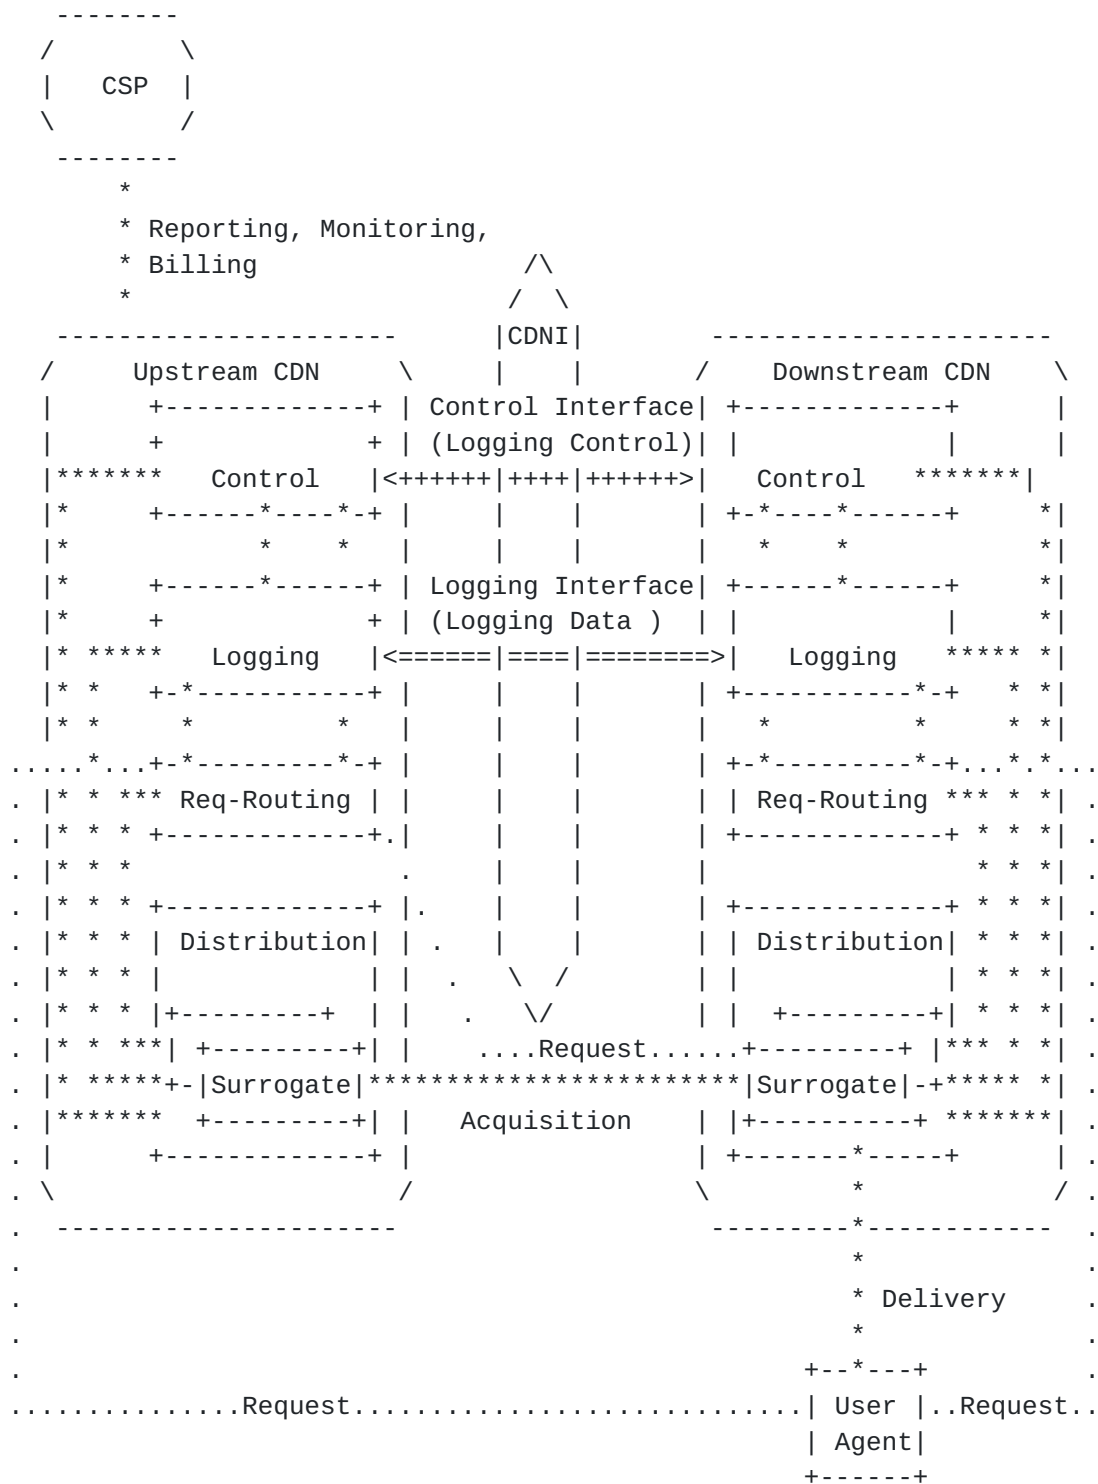
Figure 1: Interactions in CDNI Logging Reference Model

A dCDN (e.g. dCDN-2) integrates the relevant logging data obtained from its dCDNs (e.g. dCDN-3) in the logging data that it provides to the uCDN, so that the uCDN ultimately obtains all logging information relevant to a CSP for which it acts as the authoritative CDN.

Note that the format of Logging data that a CDN provides over the CDNI interface might be different from the one that the CDN uses internally. In this case, the CDN needs to reformat the Logging data before it provides this data to the other CDN over the CDNI Logging interface. Similarly, a CDN might reformat the Logging data that it receives over the CDNI Logging interface before injecting it into its log-consuming applications or before providing some of this logging information to the CSP. Such reformatting operations introduce

latency in the logging distribution chain and introduce a processing burden. Therefore, there are benefits in specifying CDNI Logging format that are as close as possible from the CDN Log formats commonly used in CDNs today.

Figure 2 maps the CDNI Logging interactions discussed above onto the CDNI Reference Model defined in [[RFC6707](#)].



<====> CDNI Logging Interface

<++++> CDNI Control Interface

```
**** interfaces outside the scope of CDNI
```

```
.... interfaces outside the scope of CDNI
```


Figure 2: Mapping of CDNI Logging interactions on the CDNI Reference Model

As illustrated in Figure 2, the Logging Control (including signaling of which logging fields are to be communicated across CDNs for a given task) occurs over the Control Interface level. The rationale for using the Control interface for Logging Control (instead of for instance using the Metadata interface) includes:

- o the Logging Control interactions typically define fairly static information for initializing and controlling the Logging interface, which matches the role of the Control Interface as described in [[I-D.ietf-cdni-framework](#)] and [[RFC6707](#)].
- o the Logging Control information (specifying the Logging information format and scope) is primarily intended to be consumed by the (typically fairly centralized) logical entity responsible for collecting intra-CDN logs, processing, filtering those and then exporting the relevant subset of logs/fields to the other CDNs.
- o the surrogates within a given CDN are typically not expected to need to be aware of the specific set of fields or set of events that have been requested by various interconnected CDNs. Rather the surrogates are likely to perform some generic logging for all services regardless of the peculiarities of every CDNI agreement. Processing (e.g. filtering, format adaptation) of the generic logging information generated by the Surrogates is expected to take place to ensure that each interconnected CDN receives the specific set of fields and logs it has requested through Logging Control. Therefore there is no need to ensure that the Logging control information be easily distributable through the CDNs right down to surrogates.
- o the Control interface is expected to support the capability to apply control at the granularity of content sets (e.g. for content Purge) which is required for Logging Control since it is expected that a CDN may require different sets of logging fields and events for different sets of content (e.g. because it only needs to perform coarse billing for a given CSP while it needs to provide detailed analytics for another CSP).

[2.2.](#) Overall Logging Chain

This section discusses the overall logging chain within and across CDNs to clarify how CDN Logging information is expected to fit in this overall chain. Figure 3 illustrates the overall logging chain within the dCDN, across CDNs using the CDNI Logging interface and

within the uCDN. For readability, the Figure only considers logging information flowing from the dCDN to the uCDN. Note that the logging chain illustrated in the Figure is obviously only indicative and varies in specific environments. For example, there may be more or less instantiations of each entity (ie there may be 4 Log consuming applications in a given CDN. As another example, there may be one instance of Rectification process per Log Consuming Application instead of a shared one.

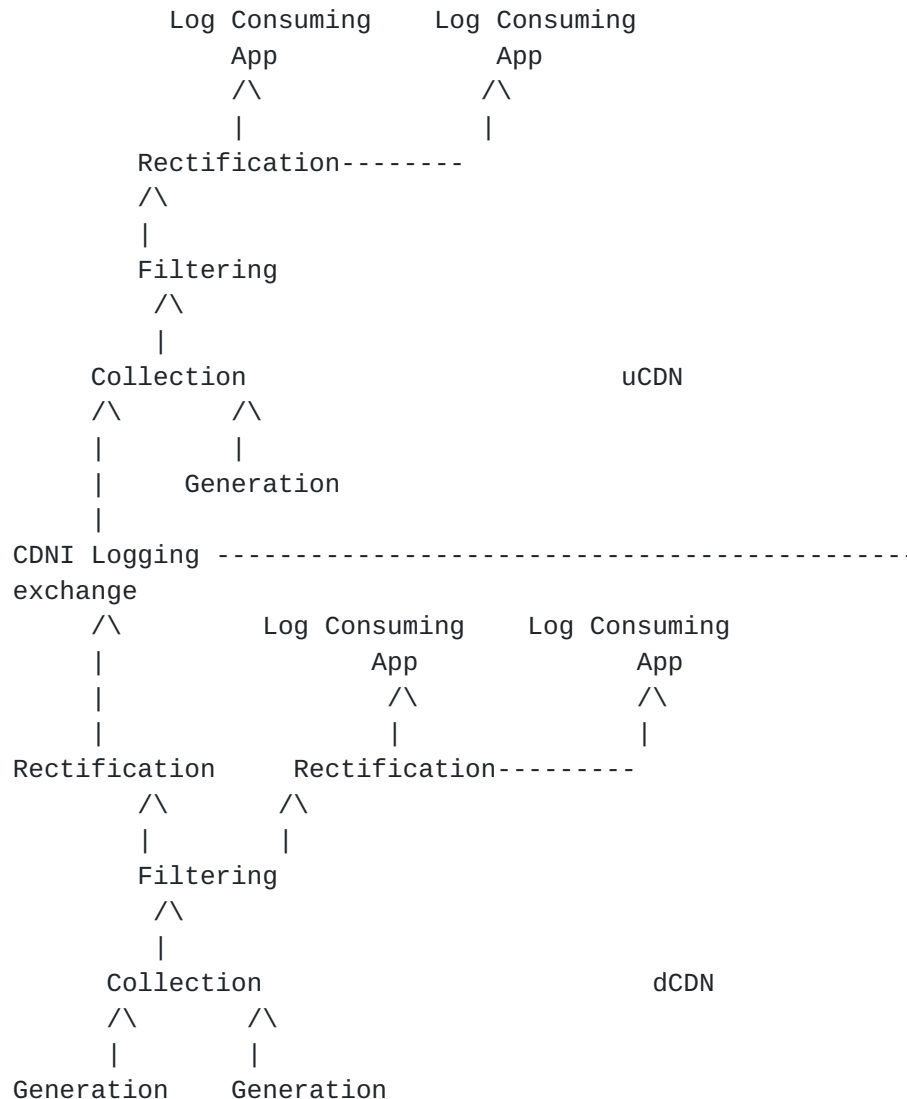


Figure 3: CDNI Logging in the overall Logging Chain

The following subsections describe each of the processes potentially involved in the logging chain of Figure 3.

2.2.1. Logging Generation and During-Generation Aggregation

CDNs typically generate logging information for all significant task completions, events, and failures. Logs are typically generated by many devices in the CDN including the surrogates, the request routing system, and the control system.

The amount of Logging information generated can be huge. Therefore, during contract negotiations, interconnected CDNs often agree on a Logging retention duration, and optionally, on a maximum size of the Logging data that the dCDN must keep. If this size is exceeded, the dCDN must alert the uCDN but may not keep more Logs for the considered time period. In addition, CDNs may aggregate logs and transmit only summaries for some categories of operations instead of the full Logging data. Note that such aggregation leads to an information loss, which may be problematic for some usages of Logging (e.g., debugging).

[I-D.brandenburg-cdni-has] discusses logging for HTTP Adaptive Streaming (HAS). In accordance with the recommendations articulated there, it is expected that a surrogate will generate separate logging information for delivery of each chunk of HAS content. This ensures that separate logging information can then be provided to interconnected CDNs over the CDNI Logging interface. Still in line with the recommendations of [I-D.brandenburg-cdni-has], the logging information for per-chunk delivery may include some information (a Content Collection Identifier and a Session Identifier as discussed in [Section 4.1.1](#)) intended to facilitate subsequent post-generation aggregation of per-chunk logs into per-session logs. Note that a CDN may also elect to generate aggregate per-session logs when performing HAS delivery, but this needs to be in addition to, and not instead of, the per-chunk delivery logs. We note that this may be revisited in future versions of this document.

2.2.2. Logging Collection

This is the process that continuously collects logs generated by the log-generating entities within a CDN.

In a CDNI environment, in addition to collecting logging information from log-generating entities within the local CDN, the Collection process also collects logging information provided by another CDN, or other CDNs, through the CDNI Logging interface. This is illustrated in Figure 3 where we see that the Collection process of the uCDN collects logging information from log-generating entities within the uCDN as well as logging information coming through CDNI Logging exchange with the dCDN through the CDNI Logging interface.

2.2.3. Logging Filtering

A CDN may require to only present different subset of the whole logging information collected to various log-consuming applications. This is achieved by the Filtering process.

In particular, the Filtering process can also filter the right subset of information that needs to be provided to a given interconnected CDN. For example, the filtering process in the dCDN can be used to ensure that only the logging information related to tasks performed on behalf of a given uCDN are made available to that uCDN (thereby filtering all the logging information related to deliveries by the dCDN of content for its own CSPs). Similarly, the Filtering process may filter or partially mask some fields, for example, to protect End Users' privacy when communicating CDNI Logging information to another CDN. Filtering of logging information prior to communication of this information to other CDNs via the CDNI Logging interface requires that the downstream CDN can recognize the set of log records that relate to each interconnected CDN.

The CDN will also filter some internal scope information such as information related to its internal alarms (security, failures, load, etc).

In some use cases described in [[I-D.ietf-cdni-use-cases](#)], the interconnected CDNs do not want to disclose details on their internal topology. The filtering process can then also filter confidential data on the dCDNs' topology (number of servers, location, etc.). In particular, information about the requests served by every Surrogate may be confidential. Therefore, the Logging information must be protected so that data such as Surrogates' hostnames is not disclosed to the uCDN. In the "Inter-Affiliates Interconnection" use case, this information may be disclosed to the uCDN because both the dCDN and the uCDN are operated by entities of the same group.

2.2.4. Logging Rectification and Post-Generation Aggregation

If Logging is generated periodically, it is important that the sessions that start in one Logging period and end in another are correctly reported. If they are reported in the starting period, then the Logging of this period will be available only after the end of the session, which delays the Logging generation.

A Logging rectification/update mechanism could be useful to reach a good trade-off between the Logging generation delay and the Logging accuracy. Depending on the selected Logging protocol(s), such mechanism may be invaluable for real time Logging, which must be provided rapidly and cannot wait for the end of operations in

progress.

In the presence of HAS, some log-consuming applications can benefit from aggregate per-session logs. For example, for analytics, per-session logs allow display of session-related trends which are much more meaningful for some types of analysis than chunk-related trends. In the case where the log-generating entities have generated during-generation aggregate logs, those can be used by the applications. In the case where aggregate logs have not been generated, the Rectification process can be extended with a Post-Generation Aggregation process that generates per-session logs from the per-chunk logs, possibly leveraging the information included in the per-chunk logs for that purpose (Content Collection IDentifier and a Session IDentifier). However, in accordance with [\[I-D.brandenburg-cdni-has\]](#), this document does not define exchange of such aggregate logs on the CDNI Logging interface. We note that this may be revisited in future versions of this document.

2.2.5. Log-Consuming Applications

2.2.5.1. Maintenance/Debugging

Logging is useful to permit the detection (and limit the risk) of content delivery failures. In particular, Logging facilitates the resolution of configuration issues.

To detect faults, Logging must enable the reporting of any CDN operation success and failure, such as request redirection, content acquisition, etc. The uCDN can summarize such information into KPIs. For instance, Logging format should allow the computation of the number of times during a given epoch that content delivery related to a specific service succeeds/fails.

Logging enables the CDN providers to identify and troubleshoot performance degradations. In particular, Logging enables the communication of traffic data (e.g., the amount of traffic that has been forwarded by a dCDN on behalf of an uCDN over a given period of time), which is particularly useful for CDN and network planning operations.

2.2.5.2. Accounting

Logging is essential for accounting, to permit inter-CDN billing and CSP billing by uCDNs. For instance, Logging enables the uCDN to check the total amount of traffic delivered by every dCDN and for every Delivery Service, as well as, the associated bandwidth usage (e.g., peak, 95th percentile), and the maximum number of simultaneous sessions over a given period of time.

2.2.5.3. Analytics and Reporting

The goal of analytics is to gather any relevant information to track audience, analyze user behavior, and monitor the performance and quality of content delivery. For instance, Logging enables the CDN providers to report on content consumption (e.g., delivered sessions per content) in a specific geographic area.

The goal of reporting is to gather any relevant information to monitor the performance and quality of content delivery and allow detection of delivery issues. For instance, reporting could track the average delivery throughput experienced by End Users in a given region for a specific CSP or content set over a period of time.

2.2.5.4. Security

The goal of security is to prevent and monitor unauthorized access, misuse, modification, and denial of access of a service. A set of information is logged for security purposes. In particular, a record of access to content is usually collected to permit the CSP to detect infringements of content delivery policies and other abnormal End User behaviors.

2.2.5.5. Legal Logging Duties

Depending on the country considered, the CDNs may have to retain specific Logging information during a legal retention period, to comply with judicial requirements.

2.2.5.6. Notions common to multiple Log Consuming Applications

2.2.5.6.1. Logging Information Views

Within a given log-consuming application, different views may be provided to different users depending on privacy, business, and scalability constraints.

For example, an analytics tool run by the uCDN can provide one view to an uCDN operator that exploits all the logging information available to the uCDN, while the tool may provide a different view to each CSP exploiting only the logging information related to the content of the given CSP.

As another example, maintenance and debugging tools may provide different views to different CDN operators, based on their operational role.

2.2.5.6.2. Key Performance Indicators (KPIs)

This section presents, for explanatory purposes, a non-exhaustive list of Key Performance Indicators (KPIs) that can be extracted/produced from logs.

Multiple log-consuming applications, such as analytics, monitoring, and maintenance applications, often compute and track such KPIs.

In a CDNI environment, depending on the situation, these KPIs may be computed by the uCDN or by the dCDN. But it is usually the uCDN that computes KPIs, because uCDN and dCDN may have different definitions of the KPIs and the computation of some KPIs requires a vision of all the deliveries performed by the uCDN and all its dCDNs.

Here is a list of important examples of KPIs:

- o Number of delivery requests received from End Users in a given region for each piece of content, during a given period of time (e.g., hour/day/week/month)
- o Percentage of delivery successes/failures among the aforementioned requests
- o Number of failures listed by failure type (e.g., HTTP error code) for requests received from End Users in a given region and for each piece of content, during a given period of time (e.g., hour/day/week/month)
- o Number and cause of premature delivery termination for End Users in a given region and for each piece of content, during a given period of time (e.g., hour/day/week/month)
- o Maximum and mean number of simultaneous sessions established by End Users in a given region, for a given Delivery Service, and during a given period of time (e.g., hour/day/week/month)
- o Volume of traffic delivered for sessions established by End Users in a given region, for a given Delivery Service, and during a given period of time (e.g., hour/day/week/month)
- o Maximum, mean, and minimum delivery throughput for sessions established by End Users in a given region, for a given Delivery Service, and during a given period of time (e.g., hour/day/week/month)
- o Cache-hit and byte-hit ratios for requests received from End Users in a given region for each piece of content, during a given period

of time (e.g., hour/day/week/month)

- o Top 10 of the most popularly requested content (during a given day/week/month),
- o Terminal type (mobile, PC, STB, if this information can be acquired from the browser type header, for example).

Additional KPIs can be computed from other sources of information than the Logging -- for instance, data collected by a content portal or by specific client-side APIs. Such KPIs are out of scope for the present memo.

The KPIs used depend strongly on the considered log-consuming application -- the CDN operator may be interested in different metrics than the CSP is. In particular, CDN operators are often interested in delivery and acquisition performance KPIs, information related to Surrogates' performance, caching information to evaluate the cache-hit ratio, information about the delivered file size to compute the volume of content delivered during peak hour, etc.

Some of the KPIs, for instance those providing an instantaneous vision of the active sessions for a given CSP's content, are useful especially if they are provided in real time. By contrast, some other KPIs, such as those averaged over a long period of time, can be provided in non-real time.

3. CDNI Logging Information Structure and Transport

As defined in [Section 1.1](#) a CDNI logging field is as an atomic logging information element and a CDNI Logging Record is a collection of CDNI Logging Fields containing all logging information corresponding to a single logging event.

This document defines non-real time transport of CDNI Logging information over the CDNI interface. For such non-real time transport, this document defines a third level of structure, the CDNI Logging File, that is a collection of CDNI Logging Records. This structure is described in Figure 4. This document then specifies how to transport such CDNI Files across interconnected CDNs. We observe that this approach can be tuned in a real deployment to achieve near-real time exchange of CDNI Logging information, e.g. by increasing the frequency of logging file creation and distribution throughout the Logging chain, but it is not expected that this approach can support real time transport (e.g. sub-second) of CDNI logging information.

+-----+									
CDNI Logging File									
+-----+									
CDNI Logging Record									
+-----+ +-----+ +-----+									
CDNI Logging CDNI Logging CDNI Logging									
Field Field Field									
+-----+ +-----+ +-----+									
+-----+									
+-----+									
CDNI Logging Record									
+-----+ +-----+ +-----+									
CDNI Logging CDNI Logging CDNI Logging									
Field Field Field									
+-----+ +-----+ +-----+									
+-----+									
+-----+									
CDNI Logging Record									
+-----+ +-----+ +-----+									
CDNI Logging CDNI Logging CDNI Logging									
Field Field Field									
+-----+ +-----+ +-----+									
+-----+									
+-----+									

Figure 4: Structure of Logging Files

It is expected that future version of this document will also specify real time transport of CDNI Logging information over the CDNI interface. We note that this might involve direct transport of CDNI Logging Records without prior grouping into a file structure to avoid the latency associated with creating and transporting such a file structure throughout the logging chain.

The semantics and encoding of the CDNI Logging fields are specified in [Section 4](#). The semantics and encoding of CDNI Records are specified in [Section 5](#). The CDNI Logging File format is specified in [Section 6](#). The protocol for transport of CDNI Logging File is specified in [Section 7](#).

4. CDNI Logging Fields

Existing CDNs Logging functions collect and consolidate logs performed by their Surrogates. Surrogates usually store the logs using a format derived from Web servers' and caching proxies' log standards such as W3C, NCSA [[ELF](#)] [[CLF](#)], or Squid format [[squid](#)]. In practice, these formats are adapted to cope with CDN specifics. [Appendix A](#) presents examples of commonly used log formats.

4.1. Generic Fields

This section specifies a set of generic CDNI Logging Fields that are expected to be found in multiple types of CDNI Logging records.

4.1.1. Semantics of Generic CDNI Logging Fields

The semantics of the generic CDNI Logging Fields are specified in Table 1.

Name	Description
Start-time	A start date and time associated with a logged event; for instance, the time at which a Surrogate received a content delivery request or the time at which an origin server received a content acquisition request.
End-time	An end date and time associated with a logged event. For instance, the time at which a Surrogate completed the handling of a content delivery request (e.g., end of delivery or error).
Duration	The duration of an operation in milliseconds. For instance, this field could be used to provide the time it took the Surrogate to send the requested file to the End-User or the time it took the Surrogate to acquire the file on a cache-miss event. In the case where Start-time, End-time, and Duration appear in a Logging Record, the Duration is to be interpreted as a total activity time related to the logged operation.
Client-IP	The IP address of the User Agent that issued the logged request or of a proxy, for instance "203.0.113.1".
Client-port	The source port of the logged request (e.g., 9542)
Destination-IP	The IP address of the host that received the logged request (e.g., 192.0.2.2).

Destination-port	The destination port of the logged request (e.g., 80).
Operation	The kind of operation that is logged; for instance, Acquisition, Delivery, or Purging.
URI_full	The full requested URL (e.g., "http://node1.peer-a.op-b.net/cdn.csp.com/movies/potter.avi?param=11&user=toto"). When HTTP request redirection is used, this URI includes the SurrogateFQDN. If the association of requests to Surrogates is confidential, the dCDN can present only URI_part to uCDN.
URI_part	The requested URL path (e.g., /cdn.csp.com/movies/potter.avi?param=11&user=toto if the full request URL was "http://node1.peer-a.op-b.net/cdn.csp.com/movies/potter.avi?param=11&user=toto"). The URI without host-name typically includes the "CDN domain" (ex.cdn.csp.com) - cf. [I-D.ietf-cdni-framework] : it enables the identification of the CSP service agreed between the CSP and the CDNP operating the uCDN.
Protocol	The protocol and protocol version of the message that triggered the Logging entry (e.g., HTTP/1.1).
Request-method	The protocol method of the request message that triggered the Logging entry.
Status	The protocol method of the reply message related to the Logging entry
Bytes-Sent	The number of bytes at application-layer protocol-level (e.g., HTTP) of the reply message related to the Logging entry. It includes the size of the response headers.
Headers-Sent	The number of bytes corresponding to response headers at application-layer protocol-level (e.g., HTTP) of the reply message related to the Logging entry.
Bytes-received	The number of bytes (headers + body) of the message that triggered the Logging entry.
Referrer	The value of the Referrer header in an HTTP request.
User-Agent	The value of the User Agent header in an HTTP request.
Cookie	The value of the Cookie header in an HTTP request.
Byte-Range	[Ed. note: to be defined]
Cache-control	The value of the cache-control header in an HTTP answer. This header is particularly important for content acquisition logs.
Record-digest	A digest of the Logging Record; it enables detecting corrupted Logging Records.

[Ed. note: specify for all Logging Fields the type (e.g., varchar, int, float, ...) and the maximum size (e.g., varchar(200))]

Name	String	Example
Time	%t	[10/Oct/2000:13:55:36-0700]
Duration	%D	-
Client-I	%a	203.0.113.45
P		
Operatio	-	-
n		
URI_full	%U	-
Protocol	%H	HTTP/1.0
Request	%m	GET
method		
Status	%>s	200
Bytes	%O	2326
Sent		
Bytes	%I	432
received		
Header	\"%{Referrer}i\"	"http://www.example.com/start.html
	\"%{User-agent}i\"	"Mozilla/4.08 [en] (Win98; I
	"	;Nav)"

Table 2: Examples using Apache format

4.2. Logging Fields for Content Delivery

Beyond the Logging Fields described in previous section, this section defines additional Logging Fields that are specifically related to Content Delivery operations. Note that the uCDN may not transfer the information provided in some of these fields to the CSP, depending on the CSP's interest in the information and on the information's confidentiality level.

4.2.1. Semantics for Delivery CDNI Logging Fields

The semantics of the generic CDNI Logging Fields are specified in Table 3.

Name	Definition
uCDN-ID	An element authenticating the operator of the uCDN as the authority having delegated the request to the dCDN.
Delivering-CDN-ID	An identifier (e.g., an aggregation of an IP address and a FQDN) of the Delivering CDN. The Delivering-CDN-ID might be considered as confidential by the dCDN. In such case, the dCDN could either not provide this field to the uCDN or overwrite the Delivering-CDN-ID with its own identifier.
Cache-bytes	The number of body bytes served from caches. This quantity permits the computation of the byte hit ratio.
Action	The Action describes how a given request was treated locally: through which transport protocol, with or without content revalidation, with a cache hit or cache miss, with fresh or stale content, and (if relevant) with which error. Example with Squid format [squid]: "TCP_REFRESH_FAIL_HIT" means that an expired copy of an object requested through TCP was in the cache. Squid attempted to make an If-Modified-Since request, but it failed. The old (stale) object was delivered to the client.

Table 3: Semantics of the Delivery CDNI Logging Fields

[Ed. note: Other information that could be logged include operations related to the authorization of the requests, URL rewriting rules enforced, the X-FORWARDED-FOR non standard HTTP header...]

4.2.2. Syntax for Delivery CDNI Logging Fields

[Ed Note: To be added]

4.3. Logging Fields for Content Acquisition

This section specifies Logging fields that are specific to Content Acquisition operations.

4.3.1. Semantics for Acquisition CDNI Logging Fields

Table 4 specifies the semantics of the Acquisition specific CDNI Logging Fields.

Name	Definition
dCDN identifier	An element authenticating the operator of the dCDN as the authority requesting the content to the uCDN
Caching_date	Date at which the delivered content was stored in cache
Validity_headers	A copy of all headers related to content validity: no-cache, ETag, Vary, last-modified...
Lookup_duration	Duration of the DNS resolution for resolving the FQDN of (uCDN's or CSP's) origin server.
Delay_to_first_bit	Duration of the operations from the sending of the content acquisition request to the reception of the first bit of the requested content.
Delay_to_last_bit	Duration of the operations from the sending of the content acquisition request to the reception of the last bit of the requested content.

Table 4: Semantics of the Acquisition CDNI Logging Fields

These information elements may be used in Content Acquisition Logging provided by dCDN to uCDN and, potentially, in Content Acquisition Logging provided by uCDN to dCDN.

4.3.2. Syntax for Acquisition CDNI Logging Fields

[Ed Note: To be added]

4.4. Logging Fields for Control

[Ed. note: LOGS RELATED TO KEY EXCHANGES FOR INSTANCE, SECTION TO BE WRITTEN AFTER THE CONTROL INTERFACE IS MORE CLEARLY DEFINED]

4.5. Logging Fields for Other Operations

Logging can be used for debugging. Therefore, all kinds of CDN operations might be logged, depending on the agreement between the dCDN and the uCDN. In particular, operations related to Request

Routing and Metadata can be logged.

5. CDNI Logging Records

[Ed. note: we need to specify the encoding of the file, the separation character, etc...]

This section defines a set of central events that a dCDN should register and publish through the Logging interface.

We classify the logged events depending on the CDN operation to which they relate: Content Delivery, Content Acquisition, Content Invalidation/Purging, etc.

5.1. Content Delivery

Some CSPs pay a lot of attention to the protection of their content (e.g., premium video CSPs). To fulfill the needs of these CSPs, a CDN shall log all the details of the content delivery authorizations. This means that a dCDN must be able to provide Logging detailing the content delivery/content acquisition authorizations and denials as well as information on why the request is authorized/denied.

CSPs and CDN service providers pay a lot of attention to errors related to content delivery. It is therefore of utmost importance that the dCDN provides detailed error information in the Logging data. This information should typically be available even when Logging is aggregated.

The content delivery events triggering the generation of a Logging Record include:

- o Reception of a content request,

The generated Logging Record typically embeds information about:

- o Denial of delivery (error or unauthorized request, e.g., HTTP 401) for a request,
- o Beginning of delivery (authorization) of a requested content,
- o End of an authorized delivery (success),
- o End of an authorized delivery (failure during the delivery, e.g., HTTP 403).

5.2. Content Acquisition

5.2.1. Logging Records Provided by dCDN to uCDN

When the uCDN requires the dCDN to provide Logging for acquisition related events, the events triggering the generation of a Logging Record include:

- o Emission of a content acquisition request (first try or retry) for a cache hit or a cache miss with content revalidation

The generated Logging Record typically embeds information about:

- o Reception of a reply indicating denial of delivery (error or unauthorized request) for a content acquisition request,
- o End of an authorized acquisition (success),
- o End of an authorized acquisition (failure)

Note that a dCDN may acquire content only from the uCDN. In this case, the uCDN can log the dCDN's content acquisition operations itself, and thus, the uCDN may not require the dCDN to log acquisition related events. However, comparing the dCDN and uCDN logs is often useful for debugging and for security auditing.

5.2.2. Logging Records Provided by uCDN to dCDN

When the dCDN requires the uCDN to provide Logging for acquisition related events, the events triggering the generation of a Logging Record include:

- o Reception of a content acquisition request for the considered Delivery Service for a cache hit or a cache miss with content revalidation

The generated Logging Record typically embeds information about:

- o Emission of a reply indicating denial of delivery (error or unauthorized request) for a content acquisition request,
- o End of an authorized acquisition (success),
- o End of an authorized acquisition (failure).

5.3. Content Invalidation and Purging

When the uCDN requests a dCDN to log invalidation/purging events (e.g., for security), the events triggering the generation of a Logging Record include:

- o Reception of a content invalidation/purging request

The generated Logging Record typically embeds information about:

- o Denial of the invalidation/purging request (error or unauthorized request, with details about the causes of the error),
- o Beginning of invalidation/purging (authorization) for a given content purging request,
- o End of an authorized invalidation/purging (success),
- o End of an authorized invalidation/purging (failure).

5.4. Logging Extensibility

Future usages might introduce the need for additional Logging fields. In addition, some use-cases such as an Inter-Affiliate Interconnection [[I-D.ietf-cdni-use-cases](#)], might take advantage of extended Logging exchanges. Therefore, it is important to permit CDNs to use additional Logging fields besides the standard ones, if they want. For instance, an "Account-name" identifying the contract enforced by the dCDN for a given request could be provided in extended fields.

The required Logging Records may depend on the considered services. For instance, static file delivery (e.g., pictures) typically does not include any delivery restrictions. By contrast, video delivery typically implies strong content delivery restrictions, as explained in [[I-D.ietf-cdni-use-cases](#)], and Logging could include information about the enforcement of these restrictions. Therefore, to ease the support of varied services as well as of future services, the Logging interface should support optional Logging Records.

6. CDNI Logging File Format

Interconnected CDNs may support various Logging formats. However, they must support at least the default Logging File format described here.

6.1. Logging Files

[Ed. Note: How many files (one per type of Delivery Service (e.g., HTTP, WMP) and per type of Event (e.g., Errors, Delivery, Acquisition,...?) and what would be inside... These aspects needs to be detailed...]

6.2. File Format

The Logging file format should be independent from the selected transport protocol, to guarantee a flexible choice of transport protocols. [Ed. note: for the real time Logging exchanges, this might be hard]

All Logging Records in a Logging File must share the same format (same set of Logging Fields, in the same order, with the same semantics, separated by the same Separator Character), to ease the parsing of the Logging data by the CDN that receives the Logging File. The CDN that provides the Logging data is responsible for guaranteeing the consistency of the Logging records' formats, typically via its log filtering and aggregation processes (see [Section 2.2.3](#)).

6.2.1. Headers

Logging files must include a header with the information described in Figure 5.

Field	Description	Examples
Format	Identification of CDNI Log format.	standard_cdni_errors_http_v1
Fields	A description of the record format (list of fields).	
Log-ID	Identifier for the CDNI Log file (facilitates detection of duplicate Logs and tracking in case of aggregation).	abcdef1234
Log-Timestamp	Time, in milliseconds, the CDNI Log was generated.	[20/Feb/2012:00:29.510+0200]
Log-Origin	Identifier of the authority (e.g., dCDN or uCDN) providing the Log- -ging	cdn1.cdni.example.com

Figure 5: Logging Headers

All time-related Logging Fields and data in the Logging File headers/footers must provide a time zone and be at least at millisecond (ms) accuracy. The accuracy must be consistent to permit the computation of KPIs involving operations realized on several CDNs.

[Ed. note: would it make sense to add a kind of "example Logging Record" in the Logging file and associated semantic (e.g. in a structure data format) ?]

6.2.2. Body (Logging Records) Format

[Ed. note: the W3C extended log format is a good base candidate to look at.]

[Ed. note: Records used for real time information and non-real time information could use different formats. In this version, we do not yet tackle the problem of real time logging exchanges]

6.2.3. Footer Format

Logging files must include a footer with the information described in Figure 6.

Field	Description	Examples
Log Digest	Digest of the complete Log (facilitates detection of Log corruption)	

Figure 6: Logging footers

This digest field permits the detection of corrupted Logging files. This can be useful, for instance, if a problem occurs on the filesystem of the dCDN Logging system and leads to a truncation of a logging file. Additional mechanisms to avoid corrupted Logging files are expected to be provided by the Logging transport protocol, cf. [Section 7](#).

7. CDNI Logging File Transport Protocol

As presented in [[RFC6707](#)], several protocols already exist that could potentially be used to exchange CDNI Logging between interconnected CDNs.

The offline exchange of non real-time Logging could rely on several protocols. In particular, the dCDN could publish the Logging on a server where the uCDN would retrieve them using a secure protocol (yet to be identified).

[Ed. note: Propose protocol, e.g. SSH File Transfer Protocol (SFTP) [[I-D.ietf-secsh-filexfer](#)]. and add call flow]

[Ed note: include options for lossless compression]

8. Logging Control

The CDNI Control interface is responsible for correctly configuring the Logging interface between interconnected CDNs, for every Delivery Service and according to the Logging configuration agreed during business negotiations.

This section will identify the parameters that the CDNI Control interface should manage on uCDN and dCDN for activating, updating, or removing a CDNI Logging configuration for a given Delivery Service.

[Ed. Note: uCDN shall be able to select the type of events that a dCDN should include in the Logging that the latter provides to the uCDN. This will be discussed during business negotiations and the Control must enforce the agreed configuration. The use of multiple levels of Logging granularity such as Syslog's "severity levels" (Emergency, Alert, Critical, ..., Debug) [[RFC5424](#)] may help in providing the most relevant amount of information depending on the intended Logging usage, as specified during the Logging format negotiation.]

[Ed. note: the specification all Logging Fields' maximum size (e.g., varchar(200)) might be constrained in some CDNs so need to exchange that information during the configuration]

9. Open Issues

The main remaining tasks on this ID are the following:

- o Detail the Logging Fields' syntax
- o Recommend a Logging File Transport Protocol and detail the call-flows
- o Detail mechanisms for Real-Time Logging

[Ed. Note: The format for Time is still to be agreed on. [RFC 5322](#) ([Section 3.3](#)) format could be used or ISO 8601 formatted date and time in UTC (same format as proposed in [[draft-caulfield-cdni-metadata-core-00](#)]). Also see [RFC5424 Section 6.2.3](#).]

[Ed. Note: When to log the end of a session when the End-User pauses a video display?]

[Ed. note: (comment from Kevin) how are errors handled ? If the client gets handed a bunch of 403s and 404s, but still gets the content eventually, without triggering an event, are those still logged? For Bytes-Sent, if there were aborted requests, do those get counted as well? Not all client behavior can be correlated with the simplified log]

10. IANA Considerations

This memo includes no request to IANA.

11. Security Considerations

11.1. Privacy

CDNs have the opportunity to collect detailed information about the downloads performed by End-Users. The provision of this information to another CDN introduces End-Users privacy protection concerns.

11.2. Non Repudiation

Logging provides the raw material for charging. It permits the dCDN to bill the uCDN for the content deliveries that the dCDN makes on behalf of the uCDN. It also permits the uCDN to bill the CSP for the content Delivery Service. Therefore, non-repudiation of Logging data is essential.

12. Acknowledgments

The authors would like to thank Sebastien Cubaud, Anne Marrec, Yannick Le Louedec, and Christian Jacquenet for detailed feedback on early versions of this document and for their input on existing Log formats.

The authors would like also to thank Fabio Costa, Yvan Massot, Renaud Edel, and Joel Favier for their input and comments.

Finally, they thank the contributors of the EU FP7 OCEAN project for valuable inputs.

13. References

13.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC5424] Gerhards, R., "The Syslog Protocol", [RFC 5424](#), March 2009.

13.2. Informative References

- [CLF] A. Luotonen, "The Common Log-file Format, W3C (work in progress)", 1995, <<http://www.w3.org/pub/WWW/Daemon/User/Config/Logging.html>>.
- [ELF] Phillip M. Hallam-Baker and Brian Behlendorf, "Extended Log File Format, W3C (work in progress), WD-logfile-960323", <<http://www.w3.org/TR/WD-logfile.html>>.
- [I-D.bertrand-cdni-experiments]
Faucheur, F. and L. Peterson, "Content Distribution Network Interconnection (CDNI) Experiments", [draft-bertrand-cdni-experiments-02](#) (work in progress), February 2012.
- [I-D.brandenburg-cdni-has]
Brandenburg, R., Deventer, O., Faucheur, F., and K. Leung, "Models for adaptive-streaming-aware CDN Interconnection", [draft-brandenburg-cdni-has-03](#) (work in progress), July 2012.
- [I-D.ietf-cdni-framework]
Peterson, L. and B. Davie, "Framework for CDN Interconnection", [draft-ietf-cdni-framework-01](#) (work in progress), July 2012.
- [I-D.ietf-cdni-requirements]
Leung, K. and Y. Lee, "Content Distribution Network Interconnection (CDNI) Requirements", [draft-ietf-cdni-requirements-03](#) (work in progress), June 2012.
- [I-D.ietf-cdni-use-cases]
Bertrand, G., Emile, S., Burbridge, T., Eardley, P., Ma, K., and G. Watson, "Use Cases for Content Delivery Network Interconnection", [draft-ietf-cdni-use-cases-10](#) (work in progress), August 2012.
- [I-D.ietf-secsh-filexfer]
Galbraith, J. and O. Saarenmaa, "SSH File Transfer Protocol", [draft-ietf-secsh-filexfer-13](#) (work in progress), July 2006.
- [RFC6707] Niven-Jenkins, B., Le Faucheur, F., and N. Bitar, "Content Distribution Network Interconnection (CDNI) Problem Statement", [RFC 6707](#), September 2012.

[apache] "Apache 2.2 log files documentation", Feb. 2012,
<<http://httpd.apache.org/docs/current/logs.html>>.

[squid] "Squid Log-Format documentation", Feb. 2012,
<<http://wiki.squid-cache.org/SquidFaq/SquidLogs>>.

Appendix A. Examples Log Format

This section provides example of log formats implemented in existing CDNs, web servers, and caching proxies.

Web servers (e.g., Apache) maintain at least one log file for logging accesses to content (the Access Log). They can typically be configured to log errors in a separate log file (the Error Log). The log formats can be specified in the server's configuration files. However, webmasters often use standard log formats to ease the log processing with available log analysis tools.

A.1. W3C Common Log File (CLF) Format

The Common Log File (CLF) format defined by the World Wide Web Consortium (W3C) working group is compatible with many log analysis tools and is supported by the main web servers (e.g., Apache) Access Logs.

According to [[CLF](#)], the common log-file format is as follows:
remotehost [rfc931](#) authuser [date] "request" status bytes.

Example (from [[apache](#)]): 127.0.0.1 - frank [10/Oct/2000:13:55:36 -0700] "GET /apache_pb.gif HTTP/1.0" 200 2326

The fields are defined as follows [[CLF](#)]:

Element	Definition
remotehost	Remote hostname (or IP number if DNS hostname is not available, or if DNSLookup is Off.
rfc931	The remote logname of the user.
authuser	The username that the user employed to authenticate himself.
[date]	Date and time of the request.
"request"	An exact copy of the request line that came from the client.
status	The status code of the HTTP reply returned to the client.

bytes	The content-length of the document transferred.	
+-----+	+-----+	+-----+

Table 5: Information elements in CLF format

A.2. W3C Extended Log File (ELF) Format

The Extended Log File (ELF) format defined by W3C extends the CLF with new fields. This format is supported by Microsoft IIS 4.0 and 5.0.

The supported fields are listed below [[ELF](#)].

+-----+	+-----+	+-----+
Element	Definition	
+-----+	+-----+	+-----+
date	Date at which transaction completed	
time	Time at which transaction completed	
time-taken	Time taken for transaction to complete in seconds	
bytes	bytes transferred	
cached	Records whether a cache hit occurred	
ip	IP address and port	
dns	DNS name	
status	Status code	
comment	Comment returned with status code	
method	Method	
uri	URI	
uri-stem	Stem portion alone of URI (omitting query)	
uri-query	Query portion alone of URI	
+-----+	+-----+	+-----+

Table 6: Information elements in ELF format

Some fields start with a prefix (e.g., "c-", "s-"), which explains which host (client/server/proxy) the field refers to.

- o Prefix Description
- o c- Client
- o s- Server
- o r- Remote
- o cs- Client to Server.
- o sc- Server to Client.

- o sr- Server to Remote Server (used by proxies)
- o rs- Remote Server to Server (used by proxies)

Example: date time s-ip cs-method cs-uri-stem cs-uri-query s-port cs-username c-ip cs(User-Agent) sc-status sc-substatus sc-win32-status time-taken

```
2011-11-23 15:22:01 x.x.x.x GET /file 80 y.y.y.y Mozilla/
5.0+(Windows;+U;+Windows+NT+6.1;+en-US;+rv:1.9.1.6)+Gecko/
20091201+Firefox/3.5.6+GTB6 200 0 0 2137
```

A.3. National Center for Supercomputing Applications (NCSA) Common Log Format

This format for Access Logs offers the following fields:

- o host [rfc931](#) date:time "request" statuscode bytes
- o x.x.x.x userfoo [10/Jan/2010:21:15:05 +0500] "GET /index.html HTTP/1.0" 200 1043

A.4. NCSA Combined Log Format

The NCSA Combined log format is an extension of the NCSA Common log format with three (optional) additional fields: the referral field, the user_agent field, and the cookie field.

- o host [rfc931](#) username date:time request statuscode bytes referrer user_agent cookie
- o Example: x.x.x.x - userfoo [21/Jan/2012:12:13:56 +0500] "GET /index.html HTTP/1.0" 200 1043 "http://www.example.com/" "Mozilla/4.05 [en] (WinNT; I)" "USERID=CustomerA;IMPID=01234"

A.5. NCSA Separate Log Format

The NCSA Separate log format refers to a log format in which the information gathered is separated into three separate files. This way, every entry in the Access Log (in the NCSA Common log format) is complemented with an entry in a Referral log and another one in an Agent log. These three records can be correlated easily thanks to the date:time value. The format of the Referral log is as follows:

- o date:time referrer
- o Example: [21/Jan/2012:12:13:56 +0500] "http://www.example.com/index.html"

The format of the Agent log is as follows:

- o date:time agent
- o [21/Jan/2012:12:13:56 +0500] "Microsoft Internet Explorer - 5.0"

A.6. Squid 2.0 Native Log Format for Access Logs

Squid [[squid](#)] is a popular piece of open-source software for transforming a Linux host into a caching proxy. Variations of Squid log format are supported by some CDNs.

Squid common access log format is as follow: time elapsed remotehost code/status bytes method URL [rfc931](#) peerstatus/peerhost type.

Squid also supports a more detailed native access log format:
Timestamp Elapsed Client Action/Code Size Method URI Ident Hierarchy/
From Content

According to Squid 2.0 documentation [[squid](#)], these fields are defined as follows:

Element	Definition
time	Unix timestamp as UTC seconds with a millisecond resolution.
duration	The elapsed time in milliseconds the transaction busied the cache.
client address	The client IP address.
bytes	The size is the amount of data delivered to the client, including headers.
request method	The request method to obtain an object.
URL	The requested URL.
rfc931	may contain the ident lookups for the requesting client (turned off by default)
hierarchy	The hierarchy information provides information on how the request was handled (forwarding it to another cache, or requesting the content to the Origin Server).
code	
type	The content type of the object as seen in the HTTP reply header.

Table 7: Information elements in Squid format

Squid also uses a "store log", which covers the objects currently kept on disk or removed ones, for debugging purposes typically.

[Appendix B](#). Requirements

[B.1](#). Additional Requirements

Section 7 of [[I-D.ietf-cdni-requirements](#)], already specifies a set of requirements for Logging (LOG-1 to LOG-16). Some security requirements also affect Logging (e.g., SEC-4).

This section is a placeholder for requirements identified in the work on logging, before they are proposed to the requirements draft authors.

Logging data is sensitive as it provides the raw material for producing bills etc. Therefore, the protocol delivering the Logging data must be reliable to avoid information loss. In addition, the protocol must scale to support the transport of large amounts of Logging data.

CDNs need to trust Logging information, thus, they want to know:

- o who issued the Logging (authentication), and
- o if the Logging has been modified by a third party (integrity).

Logging also contains confidential data, and therefore, it should be protected from eavesdropping.

All these needs translate into security requirements on both the Logging data format and on the Logging protocol.

Finally, this protocol must comply with the requirements identified in [[I-D.ietf-cdni-requirements](#)].

[Ed. note: cf. requirements draft: "SEC-4 [MED] The CDNI solution should be able to ensure that the Downstream CDN cannot spoof a transaction log attempting to appear as if it corresponds to a request redirected by a given Upstream CDN when that request has not been redirected by this Upstream CDN. This ensures non-repudiation by the Upstream CDN of transaction logs generated by the Downstream CDN for deliveries performed by the Downstream CDN on behalf of the Upstream CDN."]

B.2. Compliance with Requirements draft

This section checks that all the identified requirements in the Requirements draft are fulfilled by this document.

[Ed. note: to be written later]

Appendix C. CDNI WG's position on candidate protocols for Logging Transport

This section will be expanded later with the position of the WG considering the alternative candidate protocols for Logging in CDNI.

[Ed. Note: in a later version, this memo will include an analysis of candidate protocols, based upon a set of (basic) requirements, such as reliable transport mode, preservation of the integrity of the information conveyed by the protocol, etc.]

C.1. CDNI WG's position on Syslog

[Ed. note: to be written later]

[Ed. note: add a few sentences to clarify why not directly use syslog... Operational reasons...]

C.2. CDNI WG's position on SNMP

As explained in [[RFC6707](#)], "SNMP traps pose scalability concerns and SNMP does not support guaranteed delivery of Traps and therefore could result in log records being lost and the consequent CoDRs and billing records for that content delivery not being produced as well as that content delivery being invisible to any analytics platforms."

Authors' Addresses

Gilles Bertrand (editor)
France Telecom - Orange
38-40 rue du General Leclerc
Issy les Moulineaux, 92130
FR

Phone: +33 1 45 29 89 46
Email: gilles.bertrand@orange.com

Stephan Emile
France Telecom - Orange
2 avenue Pierre Marzin
Lannion F-22307
France

Email: emile.stephan@orange.com

Roy Peterkofsky
Skytide, Inc.
One Kaiser Plaza, Suite 785
Oakland CA 94612
USA

Phone: +01 510 250 4284
Email: roy@skytide.com

Francois Le Faucheur
Cisco Systems
Greenside, 400 Avenue de Roumanille
Sophia Antipolis 06410
FR

Phone: +33 4 97 23 26 19
Email: flefauch@cisco.com

Pawel Grochocki
Orange Polska
ul. Obrzezna 7
Warsaw 02-691
Poland

Email: pawel.grochocki@orange.com

